

Linear stability of flow in a differentially heated cavity via large-scale eigenvalue calculations

Elizabeth A. Burroughs ^{a,*}, Richard B. Lehoucq ^b,
Louis A. Romero ^b, Andrew G. Salinger ^c

^a*Department of Mathematics and Statistics, University of New Mexico,
Albuquerque, NM 87131*

^b*Computational Mathematics and Algorithms Department, Sandia National
Laboratories, Albuquerque, NM 87185 ¹*

^c*Parallel Computational Sciences Department, Sandia National Laboratories,
Albuquerque, NM 87185 ¹*

Abstract

We locate the onset of oscillatory instability in the flow in a differentially heated cavity of aspect ratio 2 by computing a steady state and analyzing the stability of the system *via* eigenvalue approximation. We discuss the choosing of parameters for the Cayley transformation so that the calculation of selected eigenvalues of the transformed system will most reliably answer the question of stability. We also present an argument that due to the symmetry of the problem, the first two unstable modes will have eigenvalues that are nearly identical, and our numerical experiments confirm this. We also locate a codimension 2 bifurcation signifying where there is a switch in the mode of initial instability. The results were obtained using a parallel finite element CFD code (MPSalsa) along with an Arnoldi-based eigensolver (ARPACK), a preconditioned Krylov method code for the necessary linear solves (Aztec), and a stability analysis library (LOCA).

Key words: stability, eigenvalues, Arnoldi, finite element, Navier-Stokes, natural convection

* Corresponding author.

Email addresses: bburroug@math.unm.edu (Elizabeth A. Burroughs),
rblehou@sandia.gov (Richard B. Lehoucq), lromero@mp.sandia.gov (Louis A. Romero), agsalin@sandia.gov (Andrew G. Salinger).

¹ Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed-Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000

1 Introduction

We locate the onset of oscillatory instability in the flow in a differentially heated cavity by computing a steady state and analyzing the stability of the system. In the problem of natural convection in a cavity we consider the flow in a box of width L and height H where the left vertical wall is held at constant temperature that is the negative of the right. The symmetry of this problem is important in considering the eigenvalues; we present analytical reasoning and eigenvalue calculations that demonstrate why there are two modes that have eigenvalues that are nearly identical. We emphasize that our method of linear stability analysis can identify this phenomenon, while a transient calculation would have difficulty predicting it.

The problem has been the subject of much research; we formulate the problem so that it is similar to a study by Paolucci and Chenoweth (1989). We use eigenvalue calculations to predict the onset of oscillations, where Paolucci and Chenoweth find oscillatory solutions using time dependent calculations. Paolucci and Chenoweth show that as the Rayleigh number is increased (based on the ΔT and L), boundary layers develop on both vertical walls, and internal waves near the corners cause the oscillations to develop. Janssen and Henkes (1995), Xin and Le Quéré (1995), Xin et al. (1997), Le Quéré and Behnia (1998), and Mayne et al. (2000, 2001) have conducted transient analyses of this problem, considering various geometries. More recently Xin and Le Quéré (2001) have conducted a linear stability analysis in a square cavity using a direct method to solve the linear systems. While our work builds upon this body of knowledge, we differ in that our method allows us to solve larger systems and arbitrary geometries.

We achieve our computations by combining a general purpose massively parallel unstructured grid finite element CFD code, MPSalsa (Shadid, 1999), with an existing Arnoldi-based eigensolver, ARPACK, (Lehoucq et al., 1998) and a parallel iterative linear solver using preconditioned Krylov methods package, AZTEC (Tuminaro et al., 1999). MPSalsa discretizes the Navier-Stokes equations and applies Newton’s method to solve for the steady state. This is in contrast to the standard approach of performing a transient calculation. While tried-and-true, this latter approach does not allow the computation of unstable steady states. Our approach does detect these unstable steady states thus allowing bifurcation analysis. Also in contrast to our approach, traditional calculations of the stability of complicated flows are done in such a way that the resulting linear systems can be solved using direct methods. Because our interest is in discretized Navier-Stokes equations in general geometries that lead to linear systems of order 10^4 – 10^7 for two and three dimensional problems, direct methods (even sparse direct methods) for the linear solves or subspace iteration for the eigensolve are not an option. We will demonstrate that paral-

lel Krylov iterative methods can be reliably used for large-scale linear stability analysis on massively parallel machines.

The problem of the flow in a differentially heated cavity both exhibits interesting physical behavior and is suited for demonstrating our eigenvalue analysis capabilities. Our Cayley transform method, as implemented in the LOCA stability analysis library (Salinger et al., 2002b), allows us to locate the onset of oscillatory instabilities; in order to locate these instabilities it is necessary to compute the eigenvalue of the system with largest real part (Meerbergen and Spence, 1997). It remains an open problem in large-scale non-symmetric eigenvalue calculations to reliably verify that the rightmost eigenvalue has been computed. Without that result, those scientists and engineers interested in computing linear stability require a variety of analysis tools; here we present a Cayley transform method that we have found effective in finding the rightmost eigenvalue when the imaginary part of that eigenvalue is large. One of the goals of this paper is to convince the reader of the reliability and applicability of this method to other problems of this type.

We have chosen to focus on the advectively dominated problem of the flow in a differentially heated cavity of aspect ratio $\frac{H}{L} = 2$. Advectively dominated flows are characterized by eigenvalues that have a large imaginary part relative to the real part. This can result in two computational difficulties. First, it can be difficult to compute the eigenvalues of the discretized system. Our choice of Cayley transform along with the use of an Arnoldi-based algorithm proves to be a reliable method to overcome this difficulty. The second difficulty is that we may need to discretize the Navier-Stokes equations on a highly resolved mesh so that the real part of the eigenvalues will approximate those of the continuous system.

The performance of our code on this problem demonstrates that we do need fine meshes to accurately compute converged real parts of the eigenvalues of interest. We will show that this is due to discretization errors, not to a failure of the eigensolver to compute the correct eigenvalues. In fact, we emphasize that the eigensolver handles with ease the large systems we are studying. Because the limitation lies in the discretization, we claim that a transient finite element code would have the same difficulty accurately computing these flows.

We validate our results by comparing the calculations to published results involving time dependent numerical calculations. In addition, we verify our results via mesh refinement for the finite element discretization and by checking the residual accuracy of our computed eigenvalues and linear systems.

Our approach is as reliable as calculations accomplished with transient methods; our approach is more efficient because we use a Krylov subspace method and use a frozen Jacobian, so we avoid the non-linear solve made at every time

step by a transient calculation. While we can not guarantee that our approach will reliably locate all instabilities because of the need to intelligently pick the parameters in the Cayley transformation, we assert that this is the same risk associated with choosing the time step and integration time when detecting instabilities through time integration. Moreover, our approach also provides qualitative information on the fluid flow not otherwise available. As we show, the information from the eigensolver can readily be used to track instabilities in parameter space and to locate higher codimension bifurcations.

We organize our paper as outlined: In Section 2 we introduce our formulation of the problem of the flow in a differentially heated cavity that provides the numerical example for our study. We also state the Navier-Stokes equations with the Boussinesq approximation governing the motion of the flow and present a novel result regarding the symmetry of the problem and the resulting nearly identical eigenvalues. In Section 3 we discuss the finite element code MPSalsa, the Cayley transform as implemented in the LOCA library, the choice of Cayley parameters and the Arnoldi-based eigenvalue package ARPACK. Section 4 gives linear stability analysis results for convection differentially heated cavity, including comparisons with published results and mesh resolution studies. In Section 5 we highlight some of the numerical issues that arise in the linear stability analysis. Section 6 presents results of tracking instabilities as a function of the Prandtl number, including the detection of a codimension 2 bifurcation. Section 7 summarizes our findings.

2 Problem formulation

In this section we describe the problem of convection in a two dimensional vertical cavity and give the basic equations that govern our flow. We present the novel result that due to the symmetry of the problem we have a pair of nearly identical eigenvalues.

2.1 The problem of flow in a differentially heated cavity

We consider the flow in a cavity of width L and height H . The left vertical wall is held at a constant temperature $-\Delta T/2$, and the right vertical wall is held at the temperature $\Delta T/2$. We impose no-flux boundary conditions at the horizontal walls and no-slip boundary conditions on all walls.

We solve the Navier-Stokes equations with the Boussinesq approximation for the flow of a thermally driven incompressible fluid:

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} + \frac{1}{\rho} \nabla p = \nu \nabla^2 \mathbf{u} + g\beta(T - T_{ref})\mathbf{e}_g \quad (1)$$

$$\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T = \kappa \nabla^2 T \quad (2)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (3)$$

where $\mathbf{u} = u\mathbf{e}_x + v\mathbf{e}_y + w\mathbf{e}_z$, p and T are the velocity, pressure and temperature; ρ , ν and κ are the density, kinematic viscosity and thermal diffusivity; g and β are the acceleration of gravity and the thermal expansion coefficient of the fluid. The vector \mathbf{e}_g is a unit vector in the direction of the gravity vector. The Boussinesq approximation assumes that the temperatures T are all close enough to an average temperature T_{ref} that we can ignore the variations in density in all terms in the equations except for the forcing term due to gravity. In these equations we subtract the hydrostatic part of the pressure.

The boundary conditions zero velocities on all four walls, adiabatic Neumann conditions on the top and bottom walls for the heat equations, and Dirichlet temperatures on the side walls:

$$T(-\frac{L}{2}, y) = \frac{\Delta T}{2}, \text{ and } T(\frac{L}{2}, y) = -\frac{\Delta T}{2}.$$

Other than the physical constants appearing in the equations, the only parameters appearing in our problem are the temperature difference ΔT , the characteristic geometrical length L , and the geometrical aspect ratio. The dimensionless parameters that result from the parameters are the Rayleigh number,

$$Ra = \frac{g\beta\Delta TL^3}{\kappa\nu},$$

and the Prandtl number,

$$Pr = \frac{\nu}{\kappa}.$$

We achieve the desired Rayleigh and Prandtl numbers by selecting $\rho = L = \Delta T = 1$, $g = Pr \times 10^1$, $\nu = Pr \times 10^{-3}$ and $\kappa = 1 \times 10^{-3}$. We then control the Rayleigh number using $Ra = \beta \times 10^7$.

2.2 Symmetry and near-degeneracy of the eigenvalues

Because the right vertical wall is held at a temperature that is the negative of the left vertical wall, the governing equations are invariant under the following symmetry transformations:

$$Rz(\mathbf{x}) = \begin{pmatrix} -T(-\mathbf{x}) \\ -\mathbf{u}(-\mathbf{x}) \\ p(-\mathbf{x}) \end{pmatrix} \quad (4)$$

where we are representing our solution in the form

$$z(\mathbf{x}) = \begin{pmatrix} T(\mathbf{x}) \\ \mathbf{u}(\mathbf{x}) \\ p(\mathbf{x}) \end{pmatrix}.$$

If $z(\mathbf{x})$ is a solution to our equations, then so is $Rz(\mathbf{x})$. However, it is not necessary that solutions to our equations satisfy $Rz(\mathbf{x}) = z(\mathbf{x})$.

We are analyzing the stability of symmetric solutions, so all eigenfunctions will either be symmetric or anti-symmetric. Any simple eigenfunction will either satisfy $R\phi(\mathbf{x}) = \phi(\mathbf{x})$ or $R\phi(\mathbf{x}) = -\phi(\mathbf{x})$. Symmetry can only be broken through a bifurcation, so that a solution that is initially symmetric will stay symmetric as we vary a parameter unless we encounter a bifurcation point.

When our system goes unstable, the internal waves will either oscillate in a symmetric manner or in an anti-symmetric manner. Physically we expect that if the walls are well separated, then the fluid on the left should be able to oscillate independently of the fluid on the right. In order for this to be so, we would have to be able to construct eigenfunctions where the fluid on the left oscillates but that on the right does not. The only way to do this is if we have multiple eigenvalues, with one symmetric eigenvector and the other anti-symmetric. This is not quite what occurs because the two sides are not completely separated, but we almost get this. Hence we have two eigenvalues that are almost identical to each other. This result is borne out in our eigenvalue calculations, presented in Section 4.

3 Methodology

In this section we discuss the numerical methods used by MPSalsa to locate steady state solutions of Equations (1)–(3), the formulation of the eigenvalue problem and our Cayley transform method, and the numerical solution of the eigenvalue problem.

3.1 Spatial discretization and the non-linear solve

A full description of the numerical methods in MPSalsa used to locate steady state solutions of Equations (1)–(3) is available in (Shadid, 1999) and the references listed therein. A brief overview is presented in this section.

A mesh of quadrilaterals for 2D problems and hexahedra for 3D problems is generated to cover the domain. Although the code allows for general unstruc-

tured meshes, the problem in this paper uses structured meshes. For parallel runs, the mesh is partitioned using the Chaco code (Hendrickson and Leland, 1995) in a way that will distribute work evenly while minimizing communication costs between processors. A Galerkin/least-squares finite element method (Hughes et al., 1989) (GLS-FEM) is used to discretize the time-invariant versions of the governing partial differential equations (1)–(3) into a set of nonlinear algebraic equations. This formulation includes a pressure stabilization term so that the velocity components, temperature and pressure fields can all be represented with equal order nodal basis functions. GLS-FEM is a consistent stabilized scheme because when the exact solution is inserted, the Boussinesq equations are satisfied exactly. We use bilinear and trilinear nodal elements for two and three dimensional problems, respectively.

Discretization of (1)–(3) results in the matrix equation

$$\begin{pmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{N} & \mathbf{0} \end{pmatrix} \begin{bmatrix} \dot{\mathbf{u}} \\ \dot{\mathbf{p}} \end{bmatrix} + \begin{pmatrix} \mathbf{K}_{u,T} + \mathbf{C}(\mathbf{u}) & -\mathbf{D} \\ \mathbf{D}^T + \mathbf{G} & \mathbf{K}_p \end{pmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} - \begin{bmatrix} \mathbf{g} \\ \mathbf{h} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix} \quad (5)$$

where \mathbf{u} is the vector of fluid velocity components and temperature unknowns, \mathbf{p} is the pressure, \mathbf{M} is the symmetric positive definite matrix of the overlaps of the finite element basis functions, $\mathbf{K}_{u,T}$ is the stiffness matrix associated with velocity and temperature, $\mathbf{C}(\mathbf{u})$ is the nonlinear convection, \mathbf{D} is the discrete (weak) gradient, \mathbf{D}^T is the discrete (weak) divergence operator and \mathbf{K}_p is the stiffness matrix for the pressure. \mathbf{G} , \mathbf{K}_p , \mathbf{N} are stabilization terms arising from the GLS-FEM. The vectors \mathbf{g} and \mathbf{h} denote terms due to boundary conditions and the Boussinesq approximation.

The resulting nonlinear algebraic equations arising from setting the time derivative terms to zero are solved using a fully coupled Newton-Raphson method (Shadid et al., 1997). An analytic Jacobian matrix for the entire system is calculated and stored in a sparse matrix storage format. At each Newton-Raphson iteration, the linear system is solved using the Aztec package (Tuminaro et al., 1999) of parallel preconditioned Krylov iterative solvers. The accuracy of the steady state solve is set by the following stopping criterion,

$$\left(\frac{1}{N} \sum_{i=1}^N \left(\frac{|\delta_i|}{\epsilon_R |x_i| + \epsilon_A} \right)^2 \right)^{\frac{1}{2}} < 1.0, \quad (6)$$

where ϵ_R and ϵ_A are the relative and absolute tolerances desired, δ_i is the update for the unknown x_i and N is the total number of unknowns. We use relative and absolute tolerances of 10^{-5} and 10^{-8} , respectively, for this study. In Aztec we exclusively use an unrestarted GMRES iteration with a non-overlapping Schwarz preconditioner where an ILU preconditioner is used on each sub-domain (each processor contains one sub-domain). These methods

enable rapid convergence to both stable and unstable steady state solutions. The scalability of these methods to large system sizes and numbers of processors is demonstrated by the solution of a 16 million unknown model on 2048 processors (Burroughs et al., 2001).

3.2 The discretized eigenvalue problem and Cayley transforms

The GLS-FEM results in a spatial discretization of the Navier-Stokes equations with the Boussinesq approximation. This leads to a finite dimensional system of differential algebraic equations of the form

$$\mathbf{B}\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x}), \quad \mathbf{x}(0) = \mathbf{x}_0, \quad (7)$$

where the matrix \mathbf{B} is singular (due to the divergence free constraint) and \mathbf{x} is a vector containing the nodal values of the velocities, temperature and pressure at the nodes of the finite element mesh. Because of the stabilization terms in the GLS discretization, \mathbf{B} , the matrix associated with the time derivative term in (5), is a non-symmetric matrix.

One can determine the stability of a steady state solution \mathbf{x}_s of $\mathbf{F}(\mathbf{x}_s) = 0$ in one of two ways: by solving the generalized eigenvalue problem that results from the linearization of (7) about the steady state, or by using a time integration scheme.

The first approach solves the generalized eigenvalue problem

$$\lambda \mathbf{B} \mathbf{z} = \mathbf{J}(\mathbf{x}_s) \mathbf{z} \equiv \mathbf{J} \mathbf{z}. \quad (8)$$

that arises from the linearization of (7) about the steady state. The matrix $\mathbf{J}(\mathbf{x}_s)$ is the Jacobian of $\mathbf{F}(\cdot)$ linearized about \mathbf{x}_s . We assume that the eigenvalues are ordered with respect to decreasing real part; $\text{real}(\lambda_{i+1}) \leq \text{real}(\lambda_i)$. If all the eigenvalues of (8) have negative real parts, the steady state is stable.

We use a Cayley transform so that we find the eigenvalues γ_i of the system

$$(\mathbf{J} - \sigma \mathbf{B})^{-1}(\mathbf{J} - \mu \mathbf{B}) \mathbf{z} = \gamma \mathbf{z}$$

that are related to the eigenvalues λ_k of (8) via

$$\gamma_i = \frac{\lambda_k - \mu}{\lambda_k - \sigma} \quad i = 1, \dots, n; k = 1, \dots, n$$

We choose $\sigma > 0$ and $\mu = -\sigma$; we choose the value of σ so that it is of similar magnitude to the imaginary part of the eigenvalue of interest, and so that

$\sigma > \text{Re}(\lambda_1)$. This transformation has the property of mapping a λ in the right half of the complex plane (i.e. an unstable mode) to a γ outside the unit circle, and those on the left halfplane (i.e. a stable mode) to a γ inside the unit circle. That is,

$$\text{real}(\lambda) > 0 \implies \|\gamma\| > 1.0, \text{ and } \text{real}(\lambda) < 0 \implies \|\gamma\| < 1.0.$$

Since Arnoldi's method will converge more rapidly to those eigenvalues with larger magnitudes, this is a very desirable property for calculating eigenvalues for use in linear stability analysis.

The use of preconditioned Krylov methods for both the eigenvalue problem and ensuing linear solves for large-scale two and three dimensional problems is not generally undertaken. The results of our paper will show that we have found success in this method. The computation of eigenvalues of the linearized steady state has received much attention in the last fifteen years (Christodoulou and Scriven, 1988; Cliffe et al., 1993; Edwards et al., 1994; Mittelman et al., 1994; Fortin et al., 1997; van Dorsselaer, 1997; Morzyński et al., 1999; Tukerman et al., 2000; Lehoucq and Salinger, 2001). The consensus of this research is to convert the generalized eigenvalue problem (8) into a standard eigenvalue problem and then solve the resulting set of linear equations during each iteration of the eigensolver. Most of the authors of these papers then solve the eigenvalue problem using inverse subspace iteration or Arnoldi's method with a sparse direct method for the resulting linear set of equations (Christodoulou and Scriven, 1988; Cliffe et al., 1993; Fortin et al., 1997; Morzyński et al., 1999; van Dorsselaer, 1997; Mittelman et al., 1994). This typically limits the linear stability analysis to two dimensional problems. It is not clear at the outset that our approach will be successful; our approach of using Cayley transforms to reduce (8) to a standard eigenvalue problem leads Fortin et al. (1997, p.1189) to state that all such "variants that we tested failed." On the contrary, we have found success with this method, and in fact find that the eigensolver performs with ease on our large (order 10^5 – 10^7) systems.

The second approach to computing the stability of a steady state is to use a time integration scheme; standard time integration schemes typically perform a nonlinear solve (due to convection) at every time step. We can think of these as computing an iteration of the form

$$\mathbf{x}^{n+1} = \mathbf{G}(\mathbf{x}^n). \tag{9}$$

The iteration is initialized with an iterate near the steady state and if the iteration converges towards the fixed point \mathbf{x}_s , then the steady state is declared stable. If \mathbf{x}_0 is an initial condition for (9), then the convergence and numerical stability of the fixed point iteration is determined by the spectral radius of

the Jacobian of $\mathbf{G}(\cdot)$. In particular, denote the eigenvalues of $\mathbf{G}_{\mathbf{x}}(\mathbf{x}_0)$ by γ_i ordered so that $|\gamma_{i+1}| \leq |\gamma_i|$.

A popular time integration scheme is given by the trapezoidal rule and results in the iteration

$$\mathbf{x}^{n+1} = \mathbf{G}(\mathbf{x}^n) = \left(\mathbf{B} - \frac{\Delta t}{2} \mathbf{J} \right)^{-1} \left(\mathbf{B} + \frac{\Delta t}{2} \mathbf{J} \right) \mathbf{x}^n \quad (10)$$

where the Jacobian is “frozen” at the steady state. The eigenvalues γ_i and λ_i are related via

$$\gamma_i = -\frac{\lambda_k + \frac{2}{\Delta t}}{\lambda_k - \frac{2}{\Delta t}} \quad i = 1, \dots, n; k = 1, \dots, n$$

and so, in principle, the eigenvalues of (8) can be determined by computing those of

$$-(\mathbf{J} - \sigma \mathbf{B})^{-1} (\mathbf{J} - \mu \mathbf{B}) \mathbf{z} \equiv \mathbf{G} \mathbf{z} = -\gamma \mathbf{z}$$

where $\mu = -\sigma = 2/\Delta t$. Note that this is the same as our choice of Cayley transform with $\mu = -\sigma = 2/\Delta t$.

The above discussion demonstrates that at a steady state, time integration and computing the eigenvalues of (8) are intimately related when a frozen Jacobian approximation is employed. We remark that although large-scale eigensolvers (subspace iteration or Arnoldi’s method) favor the computation of those eigenvalues largest in magnitude, these may not be the desired rightmost eigenvalues. This occurs when the flow is advectively dominated. Our choice of a Cayley transform allows us to overcome this difficulty.

We now explain why Arnoldi’s method for the eigenvalue solvers is preferred to the typically undertaken transient calculation. A transient calculation (with the linearized Jacobian \mathbf{J}) or fixed point iteration is equivalent to the power method on \mathbf{G} . The rate of convergence to the eigenvector associated with γ_1 is $|\gamma_2/\gamma_1|$. The rate of convergence improves to $|\gamma_{m+1}/\gamma_1|$ if the power method is replaced by subspace iteration on m vectors. However, the resulting rate of convergence can be intolerable. The rate of convergence to $\gamma_1, \gamma_2, \dots, \gamma_r$ may be dramatically improved by projecting \mathbf{G} onto the column space of

$$\mathbf{x}^0, \mathbf{x}^1, \dots, \mathbf{x}^m.$$

Arnoldi’s method (Arnoldi, 1951) iteratively determines an orthogonal basis for the above column space that by definition is a Krylov subspace.

3.3 Arnoldi's method and the numerical solution of the eigenvalue problem

The remainder of the section reviews several issues with the use of Arnoldi's method for the numerical solution of the eigenvalue problem. We use the parallel implementation P_ARPACK (Maschhoff and Sorensen, 1996) of ARPACK (Lehoucq et al., 1998) for computing the eigenvalues of (8) via Cayley transforms. We refer the reader to (Lehoucq and Salinger, 2001) for information regarding the use of ARPACK for problems in linear stability analysis.

We discuss the selection of the Cayley parameters σ and μ . There are two strategies by which we can choose the Cayley parameters. The first strategy was presented in the previous subsection and draws upon a connection with the trapezoidal rule in fixed point iteration. This is the strategy we employ in this study; we will discuss the implications of this choice in Section 5. The second strategy was presented by Lehoucq and Salinger (2001); the Cayley parameters are selected $\lambda_1 < \sigma < \mu$ so that the condition number of $(\mathbf{J} - \sigma\mathbf{B})^{-1}(\mathbf{J} - \mu\mathbf{B})$ is bounded and so can be efficiently solved with preconditioned Krylov methods. This second strategy tends to be more efficient than the first strategy for finding eigenvalues with zero or small imaginary parts; however, it is not as reliable. (Nor is there a relationship with a common fixed point iteration scheme for determining the stability of the steady state. The analogous time-stepper is unconditionally unstable for all modes.) The lack of reliability manifests itself when the flow is advectively dominated so that the rightmost λ 's do not correspond to the largest in magnitude γ 's. We remark that we encountered this unreliability in the solution of the problem of the secondary bifurcation from steady rolls into oscillatory rolls in the Rayleigh-Bénard problem, discussed in Burroughs et al. (2001): the first strategy finds the eigenvalues of interest where the second does not.

We briefly overview several salient issues. Further details are available in the discussion of the numerical experiments performed in Section 4 and Section 5, and in the paper (Lehoucq and Salinger, 2001).

- (1) The numerical solution of the linear system resulting from using a Cayley transform is found by exclusively using an unrestarted GMRES iteration with a non-overlapping Schwarz preconditioner where an ILU preconditioner is used on each sub-domain (each processor contains one sub-domain).
- (2) We must choose the size of the Arnoldi space m (needed by ARPACK). Our findings, in general, are that for the most difficult problems m was never larger than 160 and 80 was typically more than adequate. We remark that although ARPACK does provide a capability to restart the Arnoldi iteration, our experiments did not use this capability. Instead,

our focus is to examine the use of preconditioned Krylov methods for linear stability analysis.

- (3) The tolerance needed by the GMRES iteration and ARPACK and their relationship was studied in (Lehoucq and Salinger, 2001), and adjusts automatically to the scaling of the problem. In general, these tolerances were no larger than 10^{-6} and no smaller 10^{-9} .
- (4) Since the Boussinesq equations (1)–(3) model an incompressible fluid, the starting vector for ARPACK is selected as $\mathbf{J}^{-1}\mathbf{B}\mathbf{w}$, where \mathbf{w} is a random vector. The resulting vector is divergence free (Meerbergen and Spence, 1997).
- (5) The P_ARPACK subroutines `pdnaupd` and `pdneupd` were modified to implement the Cayley transform and an improved check for termination. The eigensolve is terminated when $\lambda_1, \lambda_2, \dots, \lambda_r$ and corresponding approximate eigenvectors for a user specified r satisfy the residual tolerance. This code is available through the LOCA library (Salinger et al., 2002b).

4 Results of convection in a differentially heated cavity

We conduct our numerical experiments at a Prandtl number of .71 and with $H = 2$ and $L = 1$ in order to compare our results to those of Paolucci and Chenoweth (1989). We validate our results through comparison to the numerical solutions of Paolucci and Chenoweth and verify our results by tracking the residual accuracy of our computed eigenvalues and linear systems and through a study of convergence as we refine the finite element mesh.

Although Paolucci and Chenoweth did not make the Boussinesq approximation in their calculations, they purposely used conditions that are well approximated by the Boussinesq approximation. In particular, $\Delta T/T_{AV} = .01$ where ΔT is the difference between the wall temperatures and T_{AV} is the average of the wall temperatures. When $A = H/L = 2$ they found an at a Rayleigh number of approximately $Ra = 3 \times 10^7$ with a dimensionless frequency of $f = 173.2$. (Because we do not make our equations dimensionless in the same way, to compare the frequencies f_{PC} reported in Paolucci and Chenoweth to the imaginary part of our computed eigenvalues we look at $\omega = f_{PC} \times 2\pi/1000$.)

We solve on quadrilateral finite element meshes with bilinear basis functions of 40×80 , 80×160 , 160×320 , 320×640 and 640×1280 . The spacing between the mesh points increases exponentially as we move away from the walls, with the points in the middle of the box having mesh spacings about 20 times as large as those near the walls.

For the finest mesh, we have 3,284,484 unknowns and solve on 256 processors

Table I

The eigenvalues for convection in a cavity with mesh 160×320 . Note that ω_1 & ω_2 are based on the frequencies reported by Paolucci and Chenoweth (1989) and are available for comparison for the two Rayleigh numbers 3.0×10^7 and 2.0×10^7

$Ra(10^7)$	ω_1	ω_2	λ_1	λ_2	λ_3
3.0	1.088		$0.3295 \pm 1.097i$	$0.3259 \pm 1.099i$	$0.0961 \pm 12.14i$
2.75			$0.2678 \pm 1.056i$	$0.2634 \pm 1.0574i$	$0.0474 \pm 11.44i$
2.5			$0.1937 \pm 1.010i$	$0.1884 \pm 1.012i$	$-0.0017 \pm 10.73i$
2.25			$0.1067 \pm 0.9584i$	$0.1001 \pm 0.9628i$	$-0.0479 \pm 9.997i$
2.0		2.316	$0.0138 \pm 0.8946i$	$0.0001 \pm 0.9081i$	$-0.0681 \pm 2.338i$
1.75			$-0.0631 \pm 0.8143i$	-0.0649	$-0.0757 \pm 2.177i$

of the Sandia-Intel TFlop machine (ASCI Red) with 333 MHz Pentium processors. On this final mesh it is somewhat difficult to achieve convergence to the steady state solution; we rely on continuation to find the steady state at the desired Rayleigh numbers. The number of GMRES solves for each eigen-solver iteration is approximately 200. The time to compute eigenvalues for the finest mesh is 6 hours for $Ra = 3.0 \times 10^7$. We set the Cayley parameters $\sigma = 5$, $\mu = -5$ and the Arnoldi size to 160.

Table I shows the eigenvalues for the 160×320 mesh and how they compare with the results of Paolucci and Chenoweth. Paolucci and Chenoweth performed calculations at Rayleigh numbers of 3×10^7 and 2×10^7 for $A = 2.0$. The frequency they report at $Ra = 3 \times 10^7$ is in excellent agreement with the frequency predicted by our eigenvalue calculation (1.088 vs. 1.097). When $Ra = 2 \times 10^7$ we still get good agreement (2.316 vs. 2.338), but the frequency they report agrees with what we calculate to be the third mode. While they report the flow as being stable, our eigenvalue calculations report that the flow is unstable because of first two modes have positive real parts. Possible explanations for why the previous work may have missed this mode include that the ungraded mesh used for this data point (generated with the computing power available 14 years ago) may not have fully resolved the flow, or that the starting point for the transient calculation did not contain a significant contribution in the direction of this instability (which is very close to being neutrally stable).

In order to see how the steady state solution converges with mesh refinement we have included Table II. This table shows the three most unstable eigenvalues and the maximum value of the x-velocity calculated with our various meshes. We are clearly getting convergence, but the convergence of the maximum x-velocity with mesh is somewhat slow and clearly is no better than the convergence with mesh of the eigenvalues.

We see slow convergence towards the real parts of the most unstable eigenvalue. (Other test problems we have studied that are not strongly advectively dominated flows show quadratic convergence rates (Burroughs et al., 2001).)

Table II

Eigenvalues and maximum computed values for the velocity in the problem of the onset of convection in a heated cavity with $Ra = 3.0 \times 10^7$ with varying mesh resolution of $N \times 2N$.

N	λ_1	λ_2	λ_3	x-velocity	coordinates
40	$0.3217 \pm 1.020i$	$0.3192 \pm 1.020i$	$-0.0778 \pm 2.856i$	0.8054	(0.119, 1.97)
80	$0.3326 \pm 1.091i$	$0.3294 \pm 1.092i$	$-0.0003 \pm 11.98i$	0.7993	(0.129, 1.97)
160	$0.3295 \pm 1.097i$	$0.3295 \pm 1.099i$	$0.0961 \pm 12.14i$	0.8032	(0.124, 1.97)
320	$0.3275 \pm 1.096i$	$0.3238 \pm 1.098i$	$0.1040 \pm 12.19i$	0.8048	(0.124, 1.97)
640	$0.3267 \pm 1.096i$	$0.3231 \pm 1.097i$	$0.1039 \pm 12.20i$	0.8052	(0.122, 1.97)

We believe that this problem demonstrates the limitations of looking for grid independence with a linear basis functions, particularly for highly advective flows. However, we note that the difficulties are with the resolution of the discretization and not in solving the eigenvalue problem. We emphasize that a transient solution is not any more reliable than the eigenvalue computations, and that in fact our eigensolver encounters no trouble in this 3 million unknown system. We also note that this problem is two dimensional; if we were trying to achieve the same resolution on a three dimensional problem, we would have billions of unknowns.

5 Numerical issues

Because we use parallel preconditioned Krylov iterative methods for the eigenvalue problem and resulting linear sets of equations, our results are obtained by specifying the values of certain adjustable parameters: we need to specify the Cayley parameters σ and μ and the size of the Arnoldi space. We briefly review our verification procedures used for our numerical experiments; as several issues have been discussed in a previous paper, which used the same CFD Code, MPSalsa, and eigensolver, P_ARPACK, but a different Cayley method, the reader is referred to (Lehoucq and Salinger, 2001) for information regarding details of linear algebra tolerances. Our main emphasis in this section is to illustrate how sensitive our results are to the Cayley parameters.

Denote by λ_c and \mathbf{z}_c the approximations to an eigenvalue and eigenvector of (8). We verify these approximations by computing the norm of the residual

$$\text{Direct Residual} = \frac{\|\mathbf{J}\mathbf{z}_c - \lambda_c \mathbf{B}\mathbf{z}_c\|}{\|\mathbf{B}\mathbf{z}_c\|}, \quad (11)$$

where $\|\cdot\|$ is the Euclidean norm of a vector. These errors only vanish when λ_c and \mathbf{z}_c are an eigenpair for (8). Note that these measures are independent of the scaling of \mathbf{z}_c .

We now discuss the Cayley parameters and the size m of the Arnoldi space used by ARPACK. These two parameters are related because if one chooses the Cayley parameters poorly, a large Arnoldi space will be required to obtain accurate eigenvalues. Our experience dictates that it is best to choose the Cayley parameters so that they are of the same order of the imaginary part of the most unstable eigenvalue. We believe that this is a reasonable assumption because the user typically has some idea of the location of the imaginary part of the most unstable eigenvalue. For example, this information is available if we are solving a problem that is a small variation of a problem that has already been solved, or if we have access to related experimental or computational results. This is a drawback to the method if there is no prior evidence regarding the size of the imaginary portion of the most unstable eigenvalue. However, this is the same issue as choosing a time step size for transient runs that is not so large as to step over oscillations, or a total time that is too small to sense the oscillations.

Table III shows the errors in the most unstable eigenvalue of the onset of convection in a differentially heated cavity as a function of the Cayley parameters and the size of the Arnoldi space. These calculations were accomplished with a 160×320 mesh and a Rayleigh number of 3.0×10^7 . We see that changing the Cayley parameters from ± 1 to ± 0.5 or ± 5 does not significantly degrade the performance of the algorithm. By the time the Cayley parameters are ± 20 we are seeing some degradation in the algorithm, but we are still getting quite good convergence after 160 iterations. Choosing the Cayley parameters too large is the same as integrating in time with too small a time step; it requires more Arnoldi iterations (time steps) to detect an oscillation. Notice also that we sometimes misidentify the most unstable eigenvalue; looking at the error, though, we see that this misidentified eigenvalue has not converged to a reasonable tolerance. In these situations increasing the size of the Arnoldi space allows us to compute the eigenvalues more accurately.

The accuracy of all of these calculations can also be limited by the accuracy to which we solve our linear systems at each Arnoldi iteration. For example, in Table III we do not get appreciably better results by using an Arnoldi space of size 160 instead of 80. In general, to improve the accuracy of our eigenvalue calculations we must either increase the size of the Arnoldi space or choose a better value for $\mu = -\sigma$, or decrease the tolerance to which we solve our linear systems. We should note, however, that these eigenvalue calculations are already highly converged. Even those with residuals near 10^{-4} instead of below 10^{-7} were getting the eigenvalues correct to 3 digits. A comparison to the mesh convergence results in Table II indicate that the limiting factor in predicting the eigenvalues to the real PDE system is more likely to be the discretization than the eigensolver.

Table III

The effect of Arnoldi size and Cayley parameters on the problem of convection in a differentially heated cavity. These results are for the most unstable eigenvalue at $Ra = 3.0 \times 10^7$ and a grid of 160×320 . In the case where the eigenvalue of interest, $0.3295 \pm 1.097i$, is not identified as the most unstable eigenvalue, we have listed both eigenvalues.

$\sigma = -\mu$	Arnoldi Size	Eigenvalue	Direct Residual
0.5	40	$0.3295 \pm 1.097i$	5.063×10^{-8}
	80	$0.3295 \pm 1.097i$	4.564×10^{-8}
	160	$0.3295 \pm 1.097i$	4.564×10^{-8}
1	40	$0.3295 \pm 1.097i$	2.597×10^{-8}
	80	$0.4774 \pm 13.03i$	8.685×10^0
		$0.3295 \pm 1.097i$	2.904×10^{-8}
	160	$0.8677 \pm 17.39i$	7.885×10^0
		$0.3295 \pm 1.097i$	2.904×10^{-8}
	200	$0.3295 \pm 1.097i$	2.904×10^{-8}
5	40	$0.5761 \pm 12.98i$	3.726×10^{-1}
		$0.3295 \pm 1.097i$	7.769×10^{-5}
	80	$0.3295 \pm 1.097i$	2.448×10^{-7}
	160	$0.3295 \pm 1.097i$	7.281×10^{-8}
20	40	$0.6094 \pm 17.35i$	3.056×10^{-1}
		$0.4343 \pm 20.11 i$	4.553×10^{-1}
		$0.2769 \pm 1.060i$	4.801×10^{-2}
	80	$0.3272 \pm 1.096i$	2.256×10^{-4}
	160	$0.3300 \pm 1.098i$	8.196×10^{-5}

6 Results of tracking Hopf bifurcations

The method of determining solution stability through eigenvalue calculations of steady solutions lends itself to the use of bifurcation tracking algorithms. In this section we show how these combined capabilities can be used to provide considerable insight into the stability picture for the model problem for flow in a differentially heated cavity of aspect ratio 2.

The results in Table I indicate that, with the 160×320 mesh, the first instability for a fluid with $Pr = 0.71$ occurs for $1.75 \times 10^7 < Ra < 2.0 \times 10^7$. Using the solution vector, eigenvector, and imaginary part of the eigenvalue at $Ra = 2.0 \times 10^7$, we invoked the Hopf bifurcation tracking algorithm in LOCA (Salinger et al., 2002b) and previously used in Salinger et al. (2002a). This algorithm uses a Newton algorithm to directly solve for the Hopf bifurcation and requires a good initial guess as supplied by the eigensolver.




Fig. 1. The tracking of four Hopf bifurcations as a function of the Prandtl number shows that the **IA** mode goes unstable at the lowest Ra until $Pr = 0.681$, at which time the **WA** mode is the first to go unstable. These modes are both shown in Figure 3. (The dotted line extensions to the **WA** and **WS** branches were added to clarify that these branches will continue to lower values of Pr , but these parts have not been calculated.)

The Hopf tracking algorithm located the first instability, which we will term **IA**, at $Ra = 1.9608 \times 10^7$ and the second, **IS**, at $Ra = 1.9997 \times 10^7$. Visualization of the eigenmodes shows that the first mode is anti-symmetric with respect to the symmetry of the equations, as shown in Equation 4, while the second is the symmetric version of the same physical mode.

We became curious about how persistent was the phenomenon that the anti-symmetric mode is the first to lose stability as a function another system parameter. We tracked the Rayleigh number where the Hopf bifurcation occurs as a function of the Prandtl number. We did not find a change in the order of instability as we increased up to $Pr = 1.3$. However, when decreasing the Prandtl number to generate the **IA** and **IS** curves in Figure 1, we found that the curves cross at $Pr = 0.6368$ and $Ra = 2.908 \times 10^7$, indicating that indeed the symmetric mode becomes more unstable then the anti-symmetric mode for Prandtl numbers in the neighborhood below that.

However, verification of this double-Hopf bifurcation with the eigensolver led to the discovery of two other complex pairs of eigenvalues with positive real parts. Further computations produced the curves labeled **WA** and **WS** in Figure 1. These modes are the anti-symmetric and symmetric versions of the wall

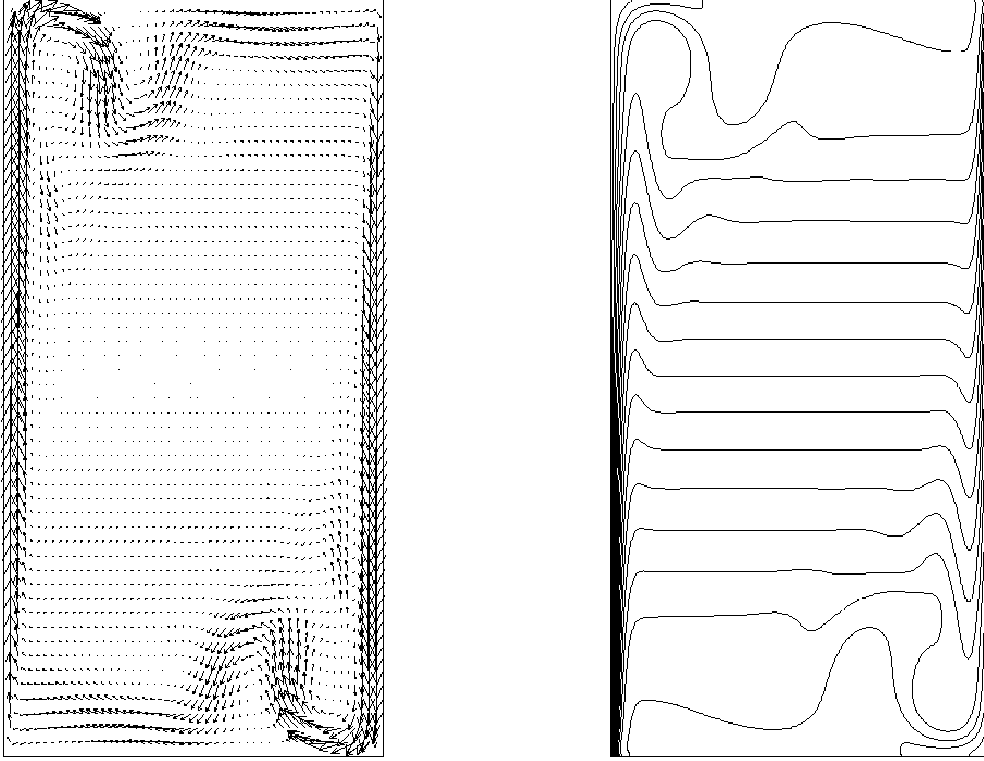


Fig. 2. This plot shows the steady solution where it loses stability with respect to two modes, at the codimension 2 bifurcation where the **IA** and **WA** branches cross in Figure 1. Velocity vectors and temperature contours are shown for this symmetric solution.

mode described by Paolucci and Chenoweth (1989). We can see graphically that a codimension 2 bifurcation occurs near $Pr = 0.681$ and $Ra = 2.22 \times 10^7$. At this Prandtl number there is a switch between whether the **IA** or **WA** mode is the first to go unstable. Figure 2 shows a visualization of the base flow and temperature contours at this point, and Figure 3 shows the temperature profiles for both modes that go unstable at this point. Since these are oscillatory instabilities, both the real and imaginary part of the eigenvectors are visualized for each mode.

As the two wall modes continue to lower Prandtl numbers, they also appear to cross. At this point, convergence was lost for the anti-symmetric mode. One interesting point is that this crossing of branches **WA** and **WS** occurs where the frequencies appear to be equal, while this was not the case when the **IA** and **IS** modes cross. This added degeneracy could be responsible for the difficulties in convergence. Since investigating this new phenomenon and its interactions with the robustness of the Hopf tracking algorithm is far outside the scope of this paper, the matter was not pursued.




Fig. 3. The modes of instability at the codimension 2 bifurcation are visualized. Temperature contours for the real and imaginary components of the anti-symmetric interior mode **IA** are shown on top, and those for the anti-symmetric wall mode **WA** are shown on the bottom.

7 Conclusions

We have completed a linear stability analysis on the problem of the flow in a differentially heated cavity. We have identified the frequency of the oscillatory instability for various Rayleigh numbers for an aspect ratio of 2 and a Prandtl number of 0.71. The frequency we identify at $Ra = 3.0 \times 10^7$ is in excellent agreement with prior published results, but for $Ra = 2.0 \times 10^7$ we find two modes more unstable than that found by Paolucci and Chenoweth (1989), and the frequency of the third most unstable mode corresponds to their published result. We also present an argument that the first two most unstable modes will have eigenvalues that are nearly identical, and our eigenvalue calculations demonstrate this is the case.

We have demonstrated both the capabilities and the limitations of using a general purpose finite element code and eigensolver for fluid stability calculations. Our interest is in large problems in possibly complex geometries where it is necessary to use iterative methods for the linear algebraic calculations. Our method has proved to be reliable in identifying the most unstable eigenvalue in advectively dominated flows because of our choice of Cayley transforms employed. The limitation of our method is that it is computationally intensive to reach high levels convergence with a low order finite element discretization. We do not believe that our eigenvalue techniques have reached any inherent limitation.

In flows that are advectively dominated, computing stability using either an eigensolver or transient calculations will produce the same difficulties in that they will require a fine discretization of the finite element mesh. We maintain that our results are as reliable as those obtained using transient integration, but that our results are more efficiently computed because we use a Krylov subspace method instead of the power method, and because we use a frozen Jacobian. We believe that our use of preconditioned Krylov iterative methods were successful because of the high quality and robust implementation of these algorithms, ARPACK and Aztec.

The determining of stability through calculation of steady states and leading eigenvalues and eigenvectors lends itself well to using bifurcation tracking algorithms. We have shown the power of using these complementary techniques by uncovering the stability behavior for a range of Prandtl number. A codimension 2 bifurcation representing the exchange of initial instability between interior and wall modes was found to exist with just a 5% decrease in the Prandtl number from the conditions previously studied.

Acknowledgements

The authors would like to acknowledge the support obtained from many members of the MPSalsa, Aztec, and LOCA teams, including David Day, Karen Devine, Sudip Dosanjh, Gary Hennigan, Scott Hutchinson, Roger Pawlowski, John Shadid, Ray Tuminaro and David Womble. We also thank Evangelos Coutsias of the University of New Mexico. This work was partially funded by the US Department of Energy through the Accelerated Strategic Computing Initiative and the Mathematical, Information and Computational Sciences programs.

References

- Arnoldi, W. E. (1951), “The principle of minimized iterations in the solution of the matrix eigenvalue problem”, *Quarterly Journal of Applied Mathematics*, Vol. 9, pp. 17–29
- Burroughs, E. A., Romero, L. A., Lehoucq, R. B. and Salinger, A. G. (2001), “Large scale eigenvalue calculations for computing the stability of buoyancy driven flows”, Technical Report SAND2001–0113, Sandia National Laboratories, Albuquerque, NM
- Christodoulou, K. N. and Scriven, L. E. (1988), “Finding leading modes of a viscous free surface flow: An asymmetric generalized eigenproblem”, *Journal of Scientific Computing*, Vol. 3, pp. 355–406
- Cliffe, K. A., Garratt, T. J. and Spence, A. (1993), “Eigenvalues of the discretized Navier-Stokes equation with application to the detection of Hopf bifurcations”, *Advances in Computational Mathematics*, Vol. 1, pp. 337–356
- Edwards, W., Tuckerman, L., Friesner, R. and Sorensen, D. (1994), “Krylov methods for the incompressible Navier-Stokes equations”, *Journal of Computational Physics*, Vol. 110, No. 1, pp. 82–102
- Fortin, A., Jarak, M., Gervais, J. and Pierre, R. (1997), “Localization of Hopf bifurcations in fluid flow problems”, *International Journal of Numerical Methods in Fluids*, Vol. 24, pp. 1185–1210
- Hendrickson, B. and Leland, R. (1995), “The Chaco user’s guide: Version 2.0”, Tech. Rep. SAND94–2692, Sandia National Labs, Albuquerque, NM
- Hughes, J. R., Franca, L. P. and Hulbert, G. M. (1989), “A new finite element formulation for computational fluid dynamics: VIII. the Galerkin/least-squares method for advective-diffusive equations”, *Computational Methods Applied Mechanics and Engineering*, Vol. 73, pp. 173–189
- Janssen, R. J. A. and Henkes, R. A. W. M. (1995), “Influence of Prandtl number on instability mechanisms and transition in a differentially heated square cavity”, *Journal of Fluid Mechanics*, Vol. 290, pp. 319–344
- Le Quéré, P. and Behnia, M. (1998), “From onset of unsteadiness to chaos in

- a differentially heated square cavity”, *Journal of Fluid Mechanics*, Vol. 359, pp. 81–107
- Lehoucq, R. and Salinger, A. (2001), “Large-scale eigenvalue calculations for stability analysis of steady flows on massively parallel computers”, *International Journal of Numerical Methods in Fluids*, Vol. 36, pp. 309–327
- Lehoucq, R. B., Sorensen, D. C. and Yang, C. (1998), *ARPACK USERS GUIDE: Solution of Large Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*, SIAM, Philadelphia, PA
- Maschhoff, K. J. and Sorensen, D. C. (1996), “P_ARPACK: An efficient portable large scale eigenvalue package for distributed memory parallel architectures”, Wasniewski, J., Dongarra, J., Madsen, K. and Olesen, D. (Eds.), “Applied Parallel Computing in Industrial Problems and Optimization”, Springer-Verlag, Berlin, Vol. 1184 of *Lecture Notes in Computer Science*
- Mayne, D. A., Usmani, A. S. and Crapper, M. (2000), “h-adaptive finite element solution of high rayleigh number thermally driven cavity problem”, *International Journal of Numerical Methods for Heat & Fluid Flow*, Vol. 10(6), pp. 598–615
- Mayne, D. A., Usmani, A. S. and Crapper, M. (2001), “h-adaptive finite element solution of unsteady thermally driven cavity problem”, *International Journal of Numerical Methods for Heat & Fluid Flow*, Vol. 11(2), pp. 172–194
- Meerbergen, K. and Spence, A. (1997), “Implicitly restarted Arnoldi with purification for the shift-invert transformation”, *Mathematics of Computation*, Vol. 218, pp. 667–689
- Mittelmann, H., Chang, K.-T., Jankowski, D. and Neitzel, G. (1994), “Iterative solution of the eigenvalue problem in Hopf bifurcation for the Boussinesq equations”, *SIAM Journal of Scientific Computing*, Vol. 15, No. 3, pp. 704–712
- Morzyński, M., Afanasiev, K. and Thiele, F. (1999), “Solution of the eigenvalue problems resulting from global non-parallel flow stability analysis”, *Computer Methods in Applied Mechanics and Engineering*, Vol. 169, pp. 161–176
- Paolucci, S. and Chenoweth, D. R. (1989), “Transition to chaos in a differentially heated vertical cavity”, *Journal of Fluid Mechanics*, Vol. 201, pp. 379–410
- Salinger, A., Lehoucq, R., Pawlowski, R. and Shadid, J. (2002a), “Computational bifurcation and stability studies of the 8:1 cavity problem”, *International Journal of Numerical Methods in Fluids*, accepted for publication
- Salinger, A. G., Bou-Rabee, N. M., Pawlowski, R. P., Wilkes, E. D., Burroughs, E. A., Lehoucq, R. B. and Romero, L. A. (2002b), “LOCA 1.0 Library of continuation algorithms: Theory and implementation manual”, Technical Report SAND2002-0396, Sandia National Laboratories, Albuquerque, NM
- Shadid, J. (1999), “A fully-coupled Newton-Krylov solution method for parallel unstructured finite element fluid flow, heat and mass transport”, *IJCFD*,

- Vol. 12, pp. 199–211
- Shadid, J., Tuminaro, R. and Walker, H. (1997), “An inexact Newton method for fully coupled solution of the Navier-Stokes equations with heat and mass transport”, *Journal of Computational Physics*, Vol. 137, pp. 155–185
- Tukerman, L. S., Bertagnolio, F., Daube, O., Le Quéré, P. and Barkley, D. (2000), “Stokes preconditioning for the inverse Arnoldi method”, Henry, D. and Bayeon, A. (Eds.), “Notes on Numerical Fluid Mechanics”, Vol. 74
- Tuminaro, R. S., Heroux, M., Hutchinson, S. A. and Shadid, J. N. (1999), “Aztec user’s guide: Version 2.1”, Technical Report SAND99-8801J, Sandia National Laboratories, Albuquerque, NM
- van Dorselaer, J. (1997), “Computing eigenvalues occurring in continuation methods with the Jacobi-Davidson QZ method”, *Journal of Computational Physics*, Vol. 138, No. CP975844, pp. 714–733
- Xin, S. and Le Quéré, P. (1995), “Direct numerical simulations of two-dimensional chaotic natural convection in a differentially heated cavity of aspect ratio 4”, *Journal of Fluid Mechanics*, Vol. 304, pp. 87–118
- Xin, S. and Le Quéré, P. (2001), “Linear stability analyses of natural convection flows in a differentially heated square cavity with conducting horizontal walls”, *Physics of Fluids*, Vol. 13(9), pp. 2529–2542
- Xin, S., Le Quéré, P. and Daube, O. (1997), “Natural convection in a differentially heated horizontal cylinder: Effects of Prandtl number on flow structure and instability”, *Physics of Fluids*, Vol. 9(4), pp. 1014–1033